Scientific Approaches and Techniques for Negotiation
A Game Theoretic and Artificial Intelligence Perspective

E.H. Gerding D.D.B. van Bragt J.A. La Poutré
CWI
P.O. Box 94079, 1090 GB Amsterdam, The Netherlands

ABSTRACT
Due to the rapid growth of electronic environments (such as the Internet) much research is currently being performed on autonomous trading mechanisms. This report contains an overview of the current literature on negotiations in the fields of game theory and artificial intelligence (AI). Game theorists have successfully developed and analyzed a variety of bargaining models in the past decades. We give an extensive overview of this theoretical work. In particular, research performed in the fields of cooperative and non-cooperative bargaining, bargaining with incomplete information, and bargaining over multiple issues is evaluated. The use and shortcomings of game-theoretical concepts in practical applications is discussed.

Simplifying assumptions frequently made in game-theoretical analyses, such as assumptions of perfect rationality and common knowledge, do not need to be made if the behavior of boundedly-rational negotiating agents is modeled directly, for instance using techniques borrowed from the field of AI. We show how different AI-techniques, such as decision trees, Q-learning, evolutionary algorithms, and Bayesian beliefs, can be used to develop a negotiation environment consisting of intelligent agents. These agents are able to adapt their negotiation strategies to changing user preferences and opponents. A survey of state-of-the art applications using AI-techniques is given in this report.

The main conclusion from this survey is that combining techniques and ideas from game theory and AI will make it possible to create robust and intelligent negotiation systems in the near future.

2000 Mathematics Subject Classification: 68T05, 91A05
1999 ACM Computing Classification System: I.2.11, J.4, I.2.6
Keywords and Phrases: game theory, artificial intelligence, bargaining, negotiation protocols, agents, learning, evolutionary algorithms
Note: Work carried out under theme SEN4 "Evolutionary Systems and Applied Algorithmics". This work has been performed in the framework of the project "Autonomous Systems of Trade Agents in E-Commerce" and has been funded by the Telematics Institute in the Netherlands.

1. Introduction

This report contains an overview of approaches and techniques concerned with bargaining. We here focus on the large body of literature that has been published in the fields of economics (in particular game theory) and artificial intelligence (AI). To give a brief impression of the rapid developments in this field, we first highlight some important breakthroughs in economic bargaining theory in this introduction. More details, extensions, and analyses can then be found in Section 2.

Perhaps surprisingly, the bargaining problem has challenged economists for decades. Yet the bargaining problem is stated very easily [35]:

> Two individuals have before them several possible contractual agreements. Both have interests in reaching agreement but their interests are not entirely identical. What "will be" the agreed contract, assuming that both parties behave rationally?

Before the path-breaking work of Nash [24] and, much later, Rubinstein [35] the bargaining problem was considered to be indeterminate. For example, in their influential work Von Neumann and Morgenstern [46] argued that the most one can say is that the agreed contract will lie in the so-called "bargaining set" (i.e., it is no worse than disagreement and there is no agreement that both parties would prefer). But because this bargaining set consists in general of an infinite number of different agreements this requirement does not yield a unique bargaining solution. A unique solution can be found, however, if the agreed contract satisfies additional axioms such as those proposed by Nash [24]. Because one can argue about which axioms are "reasonable" and which are not, Nash suggested to complement this axiomatic approach with a strategic game. This route was followed by Rubinstein [35] who proved that an important bargaining game (the "alternating-offers" game) has a unique solution. Binmore [5] then connected the fields of axiomatic and strategic bargaining by proving that the solution of Rubinstein's bargaining model coincides with the Nash bargaining solution under special circumstances.

Binmore et al. observe in a more recent paper [6] that researchers from outside the economic community are becoming more and more interested in game-theoretic work on bargaining. Interest in bargaining is especially surging in the artificial intelligence community, see the overview in Section 3 of this report.

Game theory frequently makes simplifying assumptions to facilitate the mathematical analysis. Common assumptions are for instance: (1) complete knowledge of the circumstances in which the game is played and (2) full rationality of the players. The first assumption implies that the rules of the game and the preferences and beliefs of the players are "common knowledge".[1] Game theorists traditionally model incomplete information by specifying a limited number of player "types". Each type is then uniquely determined by a set of preferences and beliefs. Players are not completely certain about the exact type of their opponent. However, the *probability* that an opponent is of a certain type is, again, common knowledge for all players. In this manner, a game of incomplete information can be transformed in a game of "imperfect" information[2] (see also Section 3.3).

The second assumption relates to the need for common knowledge on how players reason. It is assumed that players maximize their expected payoffs given their beliefs. Players have infinite computational capacity to pursue statements like "if I think that he thinks that I think..." ad infinitum. Furthermore, players are assumed to have a perfect memory.[3] These assumptions limit the practical applicability of game-theoretic results. In the field of AI, however, assumptions like complete knowledge or full rationality are not necessary because the behavior of individual agents can be modeled directly. This gives the AI approach an important advantage over more rigorous (but at the same time more simplified) game-theoretical models.

Researchers in the field of AI are currently developing software agents which should be able (in the near future) to negotiate in an intelligent way on behalf of their users. A survey of the potential of automated negotiation is given in [47, Ch. 9]. A well-known example is the agent-based heating system of the Xerox company. In this climate control system each agent controls an office thermostat and the allocation of resources is market based. Another example of negotiating agents is given in [7]. This paper describes a system in which a utility agent (acting on behalf of an electricity company) is negotiating with consumer agents to prevent excessive peaks in the demand for electricity.

An important restriction of the above systems is that they are closed and that agents behave and interact in a predetermined way. That is, they typically consist of a "fixed" collection of "inflexible" agents. In future applications for e-commerce, multi-agent systems will need to be much more open-ended and dynamic, especially for trading, brokering, and profiling applications. In particular, it is important for the negotiating agents to be able to adapt their strategies to deal with changing opponents, changing topics and concerns, and changing user preferences. This should lead to much more advanced and universal systems.

[1] Common knowledge means that the players know what the other players know, etc., in an infinite regress.

[2] A game is said to have perfect information if (i) there are no simultaneous moves and (ii) at each decision point it is known which choices have previously been made [44, Ch. 1].

[3] Lately, much research in game theory focuses on the field of "bounded" rationality, in which players have limited computational power and/or limited hindsight. An overview of recent work in this field can be found in [37]. Binmore also gives a short discussion of this topic in [3, pp. 478-488].

Nevertheless, due to this rapidly increasing complexity, the connection between the AI approach and a game-theoretic analysis remains important. Game theory may aid in the difficult task of choosing a suitable bargaining protocol [6]. Tools and techniques from AI can be used to develop software applications and bargaining protocols which are currently beyond the reach of classical game theory.

2. Game-theoretic approaches to bargaining

Traditionally, game theory can be divided into two branches: cooperative and non-cooperative game theory. Cooperative game theory abstracts away from specific rules of a game and is mainly concerned with finding a solution given a set of possible outcomes. A topic like coalition forming is typically analyzed using cooperative game theory. Often, in real life, companies can gain profits by working together, for example by securing a larger market share or by reducing direct competition with the competitors. In such games, a surplus is created when two or more players cooperate and form a coalition. Cooperative game theory can then determine how the surplus is to be divided, given a coalition and a set of assumptions (called "axioms"). Likewise, cooperative bargaining theory determines how the surplus is to be divided which results from an agreement.

Non-cooperative game theory, on the other hand, is concerned with specific games with a well defined set of rules and game strategies. All strategies and rules are known beforehand by the players. Non-cooperative game theory uses the notion of an equilibrium strategy to determine "rational" outcomes of a game. Numerous equilibrium concepts (and subsequent refinements) have been proposed in the literature (see [44] for an overview). Some widely-used concepts are "dominant" strategies, "Nash" equilibria and "subgame perfect" equilibria. A dominant strategy is optimal in all circumstances, that is, no matter what the strategies of the other players are. This is obviously a very strong notion of an equilibrium strategy. A slightly weaker, but still very powerful, equilibrium concept is the so-called Nash equilibrium [25, 26]. The strategies chosen by all players are said to be in Nash equilibrium if no player can benefit by unilaterally changing his strategy. Nash proved that every finite game has at least one equilibrium point (in pure or mixed strategies) [25, 26]. A important refinement of a Nash equilibrium for extensive-form games (i.e., games with a tree structure) is Selten's subgame-perfect equilibrium [39, 40]. In subgame-perfect equilibrium the strategies for *each* subgame of the game tree constitute a Nash equilibrium.

An overview of bargaining literature from the field of cooperative game theory will be given in Section 2.1. In Section 2.2 several non-cooperative bargaining games are discussed. Particular attention is paid to the most important bargaining protocol: the "alternating-offers" game. In Section 2.2 bargaining over a single issue is assumed. Section 2.3 covers work on multiple-issue negotiations.

As we mentioned before, traditional game theory assumes complete information, implying that the player's preferences and beliefs are common knowledge. However, lately many researchers in game theory have focussed on the consequences of players having private information. Among other things, incomplete information could explain why inefficient deals are reached or why no deal is reached at all. For instance, the occurrence of strikes and bargaining impasses, but also the occurrence of delays in negotiations can theoretically be addressed when complete information is no longer assumed. Literature related to this topic is discussed in Section 2.4.

2.1 Cooperative bargaining theory

Cooperative game theory considers the space of possible outcomes of a game, without specifying the game itself in detail. In case of bargaining, the outcomes are often denoted in terms of utilities [3]. In case of two-player games, the outcomes are then represented by utility pairs. Cooperative bargaining theory is concerned with the question of which outcome will eventually prevail, given the set of all possible utility pairs. A particular set of possible outcomes is also referred to as a "bargaining problem".

A function which maps a bargaining problem to a single outcome is called a "solution concept". Usually, a solution concept is only valid for a certain subset of all possible bargaining problems. For instance, the first and most famous solution concept, the Nash bargaining solution [24] only applies to convex and compact bargaining sets (see also [3, pp. 180–181]). Only if these requirements are satisfied the bargaining problem can properly be called a Nash bargaining problem.

An alternative bargaining solution has been proposed by Kalai and Smorodinsky [15]. Their approach is discussed below. Both the Nash and the Kalai and Smorodinsky bargaining solutions are invariant with respect to the calibration of the players' utility scales. The "utilitarian" solution concept differs in that respect and does actually depend on how the functions are scaled. For this reason, its application is limited to those situations where inter-personal utility comparison makes any sense. Cooperative theories of bargaining are discussed in more detail in [34].

The Nash bargaining solution Nash proposed four properties, now called the "Nash axioms", which should be satisfied by rational bargainers [24],[3, p. 184]:

1. The final outcome should not depend on how the players' utility scales are calibrated. This means the following. A utility function specifies a player's preferences. However, different utility functions can be used to model the same preferences. Specifically, any strictly increasing affine transformation of a utility function models the same preferences as the original function, and should therefore yield the same outcome.

2. The agreed payoff pair should always be individually rational[4] and Pareto-efficient[5].

3. The outcome should be independent of irrelevant alternatives. Stated otherwise, if the players sometimes agree on the utility pair s when t is also a feasible agreement[6], they never agree on t when s is a feasible agreement.

4. In symmetric situations, both players get the same.

The solution which satisfies these four properties is characterized by the payoff pair $s = (x_1, x_2)$ which maximizes the so-called Nash product $(x_1 - d_1)(x_2 - d_2)$, where d_1 and d_2 are player 1's and player 2's outcomes in case of a disagreement. Nash proved that this is the *only* solution which satisfies all four axioms [24]. Given a Nash bargaining problem where the set of individually rational agreements is not empty, the Nash bargaining solution then leads to a unique outcome. Figure 1 illustrates how to construct the Nash bargaining solution for a given bargaining problem.

Due to the fourth axiom, both players are treated symmetrically if the bargaining problem is symmetric as well. In other words, if the players' labels are reversed, each one will still receive the same payoff. A more general solution attributes so-called "bargaining powers" α and β to player 1 and player 2, respectively. In this generalized or asymmetric Nash bargaining solution, the fourth axiom is abandoned and the bargaining solution comes to depend on the bargaining powers of the two players.[7] The generalized Nash bargaining solution corresponding to the bargaining powers α and β can be characterized as above as the pair s which maximizes the product $(x_1 - d_1)^\alpha (x_2 - d_2)^\beta$ [3, p. 189].

The Kalai-Smorodinsky bargaining solution The third of the Nash axioms (independence of irrelevant alternatives) has been the source of great controversy (follow the discussion in [18]). Kalai and Smorodinsky therefore proposed an alternative to this axiom, which they refer to as the "axiom of monotonicity" [15]. For a set S of individually-rational and Pareto-efficient points, let $m_i(S) = max\{s_i \mid s \in S\}$ be player i's maximum feasible utility, for $i = 1, 2$. The Kalai-Smorodinsky solution

[4] An agreement is individually rational if it assigns each player a utility that is at least as large as a player can guarantee for himself in the absence of an agreement [3, p. 178].

[5] An agreement is Pareto-efficient if no player can gain without causing a loss for the other player [3, p. 177].

[6] That is, within the set of possible agreements.

[7] What these bargaining powers represent depends on the actual (non-cooperative) game played. For example, in case of negotiating companies the bargaining powers could be determined by the strength of their respective market positions. It should be clear however, that the bargaining powers have nothing to do with the bargaining skills of the players, since perfect rationality is assumed.

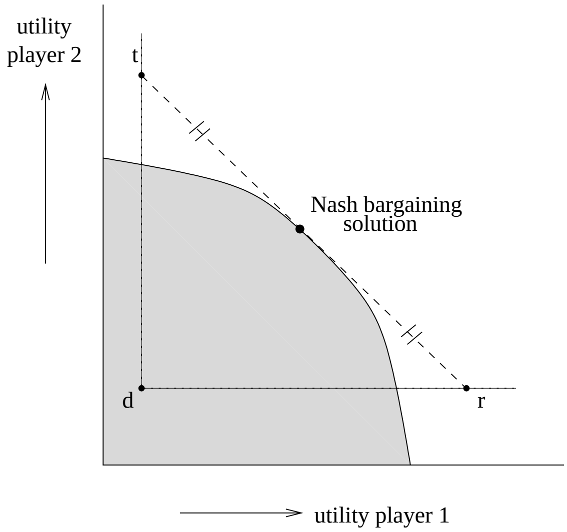

Figure 1: Construction of the Nash bargaining solution. This figure shows the Pareto-efficient frontier (denoted by the solid line) and the Nash bargaining solution for a specific bargaining problem. The bargaining problem is defined by the set of feasible utility pairs (denoted by the grey area) and the disagreement point d which specifies the players' payoffs in case of a disagreement. To find the (symmetric) Nash bargaining solution, one needs to draw a supporting line on the Pareto-efficient frontier such that the Nash bargaining solution is halfway between the points r and t. The points r and t are located on respectively the horizontal and the vertical lines drawn from the disagreement point d.

then selects the maximum element in S on the line that joins the disagreement point (d_1, d_2) with the point $(m_1(S), m_2(S))$. An example is given in figure 2.

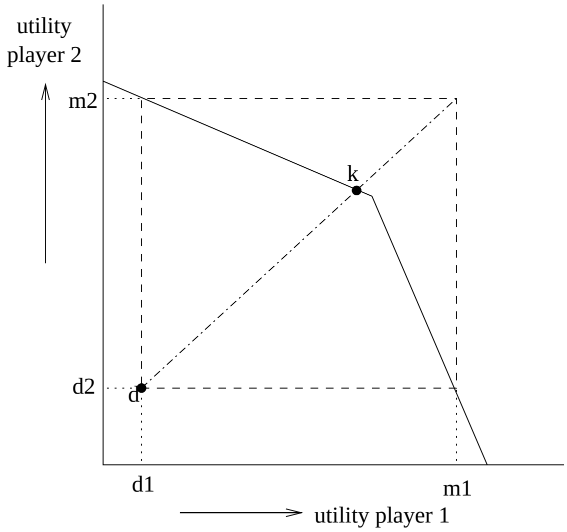

Figure 2: Construction of the Kalai-Smorodinsky solution. m_1 and m_2 are the maximum feasible utilities for players 1 and 2, respectively. Point k is the unique solution which satisfies the four axioms proposed by Kalai and Smorodinsky [15].

Utilitarianism A utilitarian policy in philosophy is one which prefers an outcome which maximizes the total welfare of the individuals in a society [23]. Any bargaining solution which maximizes the sum of utilities is therefore called a utilitarian solution concept. Stated less formally, the utilitarian principle asserts that "you should do something for me if it will hurt you less than it will help me". Clearly, a utilitarian solution concept assumes that interpersonal utility comparisons are possible. Therefore, Nash's first axiom (independence of utility calibrations) no longer holds in utilitarian models.[8]

Conclusions Apparently, many different types of solutions to the bargaining problem exist in cooperative game theory. The choice of a specific solution is of course based on norms existing in a society, or, more specifically, on which axioms seem to be "reasonable" in a specific bargaining context. Certain outcomes might be for instance be considered as "unfair". An example is given in [32, pp. 235–250].

Additionally, it is important to consider for which classes of non-cooperative games the solution concepts from cooperative game theory are appropriate. For instance, if no non-cooperative game can be found which results in a solution specified by cooperative game theory, then the results from cooperative game theory have little bearing. Fortunately, such a connection between cooperative and non-cooperative game theory has been observed under special circumstances [5] (more details are given in Section 2.2).

2.2 Bargaining over a single issue

Four different negotiation games or "protocols" are described in this section. These protocols can be used by two bargainers to divide a given bargaining "surplus", that is, the total profit resulting when the players reach an agreement. Without loss of generality, we assume that the bargaining surplus is of size unity in the remainder of this report.

The following protocols are considered below: (1) the Nash demand game, (2) the ultimatum game, (3) the alternating-offers game and (4) the monotonic concession protocol. The first three games are well-known and widely-used. The fourth game is described in [33] and is an attempt to model a more realistic negotiation scenario. However, in all games described here analytical solutions are obtained using the strong assumption of common knowledge. The extrapolation of results obtained here to real-world cases is therefore a non-trivial step.

The protocols described in this section have been applied mainly to evaluate negotiations over a single issue. In real life, this issue is often the price of a good to be negotiated. Although this keeps matters simple, important value-added services such as delivery time, warranty or service are left out. Both the supplier and the consumer could for instance benefit if negotiations involve multiple issues. Moreover, multiple-issue negotiations can be less competitive because solutions can be sought which satisfy both parties. Multiple-issue negotiations are studied in more detail in Section 2.3.

The Nash demand game Both players simultaneously demand a certain fraction of the bargaining surplus in this game, without any knowledge of the other player's demand [3, pp. 299-304]. In case the sum of demands exceeds the surplus, both players only receive their disagreement payoff. Otherwise, the demands are said to be compatible, and both players get what they requested. This game has an *infinite* number of Nash equilibria: all deals which are Pareto-efficient, but also deals where both players receive their disagreement payoff. For example, if both players ask more than the entire surplus, no player could ever gain by unilaterally changing his strategy.

The concept of a Nash equilibrium thus places few restrictions on the nature of the outcome. Nash therefore suggested a refinement for this game which does result in a *unique* solution. This refinement of the demand game is called the "perturbed" demand game [29, pp. 77-81]. In this perturbed game the players are completely certain about which outcomes are within the bargaining set (i.e., the set of compatible demands) and which outcomes are not. When the degree of uncertainty approaches

[8] Note that the Pareto-efficiency axiom still holds. The other axioms depend on the specific solution concept.

zero, the Nash equilibrium of the perturbed game approaches the Nash bargaining solution of the regular demand game (without uncertainty).[9] The reader is referred to [44] for technical details on this subject. A more introductory overview is given by Binmore [3].

The ultimatum game Playing Nash's demand game, both players could easily receive nothing, or it could occur that some of the surplus is "thrown away". Players would do better by choosing a somewhat less competitive game. If they are unable to reach an agreement using this alternative game, the demand game still remains an option.

A very simple alternative is the so-called "ultimatum" game. In this game, one of the players proposes a split of the surplus and the other player has only two options: accept or refuse. In case of a refusal, both players get nothing (or the demand game is played). Although the game again has an infinite number of Nash equilibria, it has only one subgame perfect equilibrium (in case the bargaining surplus can be divided with arbitrary precision) where the first player demands the whole surplus and the second player accepts this deal [3, pp. 197-200].

The alternating-offers game Basically a multiple-stage extension of the ultimatum game, the "alternating-offers" game is probably the most elegant bargaining model. As in the ultimatum game, player 1 starts by offering a fraction x of the surplus to player 2. If player 2 accepts player 1's offer, he receives x and player 1 receives $1 - x$. Otherwise, player 2 needs to make a counter offer in the next round, which player 1 then accepts or rejects (sending the game to the next round). This process is repeated until one of the players agrees or until a finite deadline is reached.

Bargaining over a single issue in an alternating fashion has been pioneered by Ingolf Ståhl [42]. A taxonomy and survey of economic literature on bargaining before 1972 is given in this reference. Ståhl analyzes bargaining games with a finite number of alternatives. Both games of finite and of infinite length are studied, but he primarily evaluates games of a finite length. Ståhl uses an assumption of "good-faith bargaining" to simplify the theoretical analysis. Good-faith bargaining prevents players from increasing their demands during play. He then identifies optimal strategies for rational players with perfect information by starting at the last stage of the game and then inductively working backwards until the beginning of play. This procedure yields those equilibria which can be found with dynamic programming methods.

A straightforward dynamic programming approach can fail in case of imperfect information [44, Ch. 1]. Sensible strategies can then be found by requiring that each player's optimal strategy for the entire game also prescribes an optimal strategy in every subgame. As mentioned before, this concept of a subgame-perfect equilibrium (SPE) is due to Selten [39, 40]. Rubinstein [35] successfully applied this equilibrium concept to identify a unique solution in his variant of the alternating-offers game. Rubinstein's game [35] has an infinite length and there is a continuum of alternatives. To simplify the analysis, Rubinstein made several assumptions with regard to the players' preferences. An important difference with Ståhl's model is that time preferences are assumed to be stationary (this means that the preferences of getting a part x of the surplus at time t over getting y at $t+1$ is independent of t).

Rubinstein analyzes two specific stationary models: one in which each player has a fixed bargaining cost for each period (c_1 and c_2) and one in which each player has a fixed discounting factor (δ_1 and δ_2). These discount factors model how impatient the player is [3, p. 202]. Player i 's utility for getting a fraction x of the surplus at time t is equal to $x(\delta_i)^t$. If the discount factor is smaller than 1, a deal is therefore worth less if the agreement is reached in the future than if a deal is reached immediately.

Using stationarity and other assumptions, Rubinstein first demonstrated that the Nash equilibrium concept is too weak to identify a unique solution by proving that every partitioning of the surplus can be supported as the outcome of Nash equilibrium play. To overcome this difficulty, Rubinstein then applied the concept of a SPE and proved that there exists a *unique* SPE in the alternating-offers bargaining model. For example, if both players have a fixed discounting factor (δ_1 and δ_2) the only

[9]Note that only the Nash equilibria which result in solutions within the bargaining set are considered. Nash equilibria in which no agreement is reached still remain [29, p.79].

SPE is one in which player 1 gets $(1-\delta_2)/(1-\delta_1\delta_2)$ and player 2 the remainder (of a surplus of size 1). Furthermore, if both players use their SPE strategy, agreement will be reached in the first round of the game. Notice that Rubinstein's proof assumes that both players have perfect information about the other player's preferences (i.e., their bargaining cost or discount factor). Bargaining with imperfect information (i.e., where uncertainty plays a crucial role) is discussed further in Section 2.4.

Rubinstein's paper has been very influential in bargaining theory. At the moment, a vast body of literature exists on infinite-horizon games. An overview is given in [29, 22]. Many pointers to the literature are given in these references. We will conclude this section by discussing a few key papers in this field.

An particularly important paper is [5]. In this paper a relation between the SPE outcome of the alternating-offers game and the Nash bargaining solution is identified in case of weak player preferences (e.g., discount factors close to unity or small time intervals between rounds). This establishes a link between non-cooperative and cooperative bargaining theory and justifies the use of the Nash bargaining solution to resolve negotiation problems (at least in case of complete information).

Van Damme et al. [45] have investigated the role of a smallest monetary unit (i.e., a finite number of alternatives) in the alternating-offers game with payoff discounting. They show that in case of a finite number of alternatives, any partition of the surplus can be supported as the result of a subgame-perfect equilibrium if the time interval between successive rounds becomes very small. This means that Rubinstein's assumption of a continuous spectrum of bids is essential in deriving a unique solution of the alternating-offers game under these conditions.

Recent theoretical work by Binmore et al. [4] examines an evolutionary variant of Rubinstein's game. In this paper the agents are modeled as (boundedly-rational) automata instead of perfectly-rational bargainers. In particular, the bargaining agents do not know the other agents' preferences. Binmore et al. characterize so-called modified evolutionary stable strategies (MESSes). A MESS modifies Maynard Smith's concept of a neutrally stable strategy [41] by favoring a more simple strategy over a more complex one in case of equal payoffs. Binmore et al. show theoretically that if both agents use a MESS, this constitutes a Nash equilibrium in which immediate agreement is reached. Furthermore, each agent's share of the surplus is bounded between the shares received by the two agents in the SPE of the infinite-horizon game studied by Rubinstein. These bounds collapse on the SPE partitions when the breakdown probability becomes very small, or when the players' discount factors (modeling their time preferences) become large.

Monotonic concession protocol A more restricted protocol, compared to the alternating-offers game, is described in [33]. In this "monotonic concession" protocol the two players announce their proposals simultaneously. If the offers of both agents match or exceed the other agent's demand, an agreement is reached. A coin is tossed to choose one of the offers in case they are dissimilar.

If no agreement is reached, the players need to make new offers in the next round. The offers need to be monotonic, that is, the players are not allowed to make offers which have a lower utility for their counter player compared to the last offer. Hence, a player can either make the same offer (to stand firm) or concede. Negotiations end if both agents stand firm in the same round. The players receive their disagreement payoffs in this case. Because each round at least one of the players has to make a concession (or a disagreement occurs), the protocol has a finite execution time if the minimum concession per round is fixed and larger than zero.

Note that in order to make a (monotonic) concession possible, a player needs to have some knowledge about the other players' preferences. This knowledge is crucial when several issues are negotiated at the same time. In this case not only the "sign" of the utility function, but also the relative importance of the issues becomes important.

Rosenschein and Zlotkin discuss which kinds of strategies are stable and efficient when using this protocol (in negotiations over a single issue). A strategy pair is efficient in this case if an agreement is always reached. Stability is defined using the notion of a symmetric Nash equilibrium.[10] Note

[10] A strategy s constitutes a symmetric Nash equilibrium if player 1 can do no better than playing s, given that player

that a strategy s in which both players make a concession in the same round is not stable: one of the players could do better by standing firm. On the other hand, a strategy where a player tosses a coin to determine whether to concede or stand still is not efficient (nor stable): a disagreement will occur with a probability of one fourth. The interested reader is referred to [33] for more details on the characteristics of this mechanism.

2.3 Bargaining over multiple issues

The above situations can be described as negotiations about how to divide a surplus. This means that the negotiations are distributive: a gain for one player always creates a loss for the other player. These kinds of negotiations are also referred to as "competitive" [13]. When more than a single issue is involved, and players attach different importance to these issues, tradeoffs become an option and negotiations may become "integrative". The latter kind of negotiations is the main topic of this section. Results from cooperative game theory are discussed first, followed by a overview of results from non-cooperative game theory.

Cooperative game theory An additive scoring system or a multi-attribute utility (MAUT) function can be used to represent the relationships or trade-offs between the issues if several issues are involved.[11] However, these methods are appropriate only if the issues are preferentially independent, that is, if the contribution of one issue is independent of the values of the other issues.

Once the preferences are mapped, for instance onto a MAUT function, the bargaining set can be determined. The main goal is again to reach a Pareto-efficient outcome. Previously introduced solution concepts such as the Nash bargaining solution or the Kalai-Smorodinsky solution can be used for this purpose. Several practical considerations (concerning for example fairness of the outcome) and some instructive real-world examples are given by Raiffa in [32].

Non-cooperative game theory Four different bargaining procedures can be distinguished for multiple-issue bargaining [31] (see figure 3). In case of "global" or simultaneous bargaining all issues are negotiated at once. The second procedure is called "separate" bargaining. In this protocol the issues are negotiated independently. The final two procedures fall under the header of sequential bargaining and are distinguished by their "rules of implementation". These rules specify when the players can start enjoying the benefits of the issues which have been agreed on.[12] Three possibilities are considered in [11]. Here, however, we will only mention the most important two. Using the so-called "independent implementation" rule, an agreement on an individual issue takes effect immediately, that is, the agreed upon issues are no longer discounted. In the "simultaneous implementation" on the other hand, the players have to wait until agreement is reached on all issues before they can enjoy the benefits of it. The time it takes to agree on the remaining issues also influences the profits gained on the already agreed upon issues.

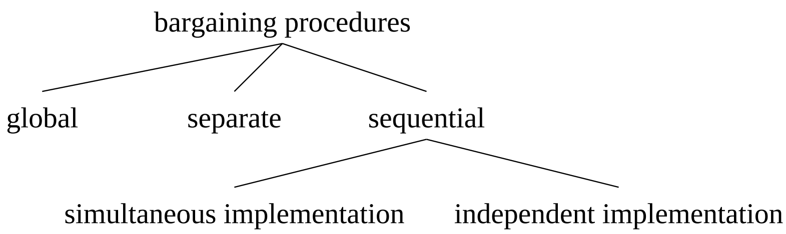

Figure 3: Four different bargaining procedures used in multiple-issue bargaining [31].

2 also uses s.

[11] See [32, pp.154-155] for a discussion of the differences between these methods.

[12] This is relevant in case the payoff is discounted in the course of time.

When bargaining is sequential an agenda needs to be determined to set the order in which the issues will be negotiated. Agenda setting is of course only relevant if the issues are of different importance. Another concern is whether the players attach the same importance to each issue or whether different players have different evaluations regarding the importance of the issues. The latter is the most interesting case since this allows for integrative negotiations. Unfortunately, however, only a limited literature exists on this topic in game theory. Usually, either the issues are of equal importance (as in [1]) or the players have identical preferences (as in [8]). In [31] the assumption is made that preferences are additive over issues, implying that the multi-issue bargaining problem is equal to the sum of the bargaining problems over the separate issues.

One of the few papers in game theory on integrative bargaining is [11]. Fershtman considers sequential bargaining over two issues. He states that, when using Rubinstein's alternating-offers protocol for each issue in a sequential order, each player prefers an agenda in which the first issue to bargain on is the one which is the least important for him but the most important for his opponent. Notably, it is shown in [11] that the subgame-perfect equilibrium outcome for this problem does not need to be Pareto-efficient.

2.4 Bargaining with incomplete information

Private information such as reservation prices (i.e. limit values on what the players find acceptable), preferences amongst issues, attitudes towards risk or time preferences are often hidden from the opponent in real-life negotiations. In bargaining, for example, it might be beneficial to be dishonest about one's attitudes towards risk in order get a greater share of the surplus (as would be the case in Rubinstein's alternating-offers game). Sometimes, however, a mechanism (or protocol) can be designed which gives agents a compelling incentive to be honest to the opponent. Such mechanisms are called "incentive compatible" and are examined in [33]. In an incentive-compatible protocol the agents can simultaneously declare their private information before the bargaining starts. The negotiations then proceed as a game of complete information.

The Vickrey auction [3, pp. 525-526] is an example of such an incentive-compatible mechanism. Unfortunately, however, only few games have this property. Therefore, it is necessary to analyze games with incomplete information. As mentioned in the introduction, games theory frequently assumes that the players have complete information. However, in order to analyze situations in which players are unsure of the opponent's type, the notion of imperfect information needs to be introduced.

Imperfect information enables us to address important issues as reputation building, signaling and self-selection mechanisms [36]. For example, the fact that players are unsure of the other player's type might explain the occurrence of (inefficient) delays in reaching an agreement [29, Ch. 5]. Using such inefficient strategies may be the only way to signal for instance one's strength (an example is the outbreak of strikes during wage bargaining situations). Any utterance which is not backed up by actions can be considered as being cheap talk.[13] Delays may therefore be required to convey private information credible [16].

In a wage negotiation problem, for example, the union is often unsure about the actual value of its workers for a firm. If this value is high, the firm will be more eager to sign an agreement. In case of a low value however, the firm will behave credible by bearing the costs of a strike [16]. A firm could try to "bluff" by ignoring a strike even in case of a high valuation, and use this strategy to signal a lower valuation of the union workers than actually is the case. However, such a strategy can potentially be very harmful.

An overview of bargaining with one-sided incomplete information is given in [29, pp. 118-120]. More introductory texts on bargaining with private information can be found in [16] and [3, Ch. 11].

[13] In non-cooperative games, nothing anyone says constrains its future behavior. If a player chooses to honor an agreement or threat that has been made, this will only be because it is optimal to do so.

3. Negotiation and learning

Several aspects of learning are potentially important during negotiation processes. First, a bargaining agent needs to have a strategy which specifies his actions during the course of play. On the basis of the agent's experiences in previous bargaining games, he can learn that it might be profitable to adjust his strategy in order to achieve better deals. Second, it might even be useful to update a strategy during play. This may be the case if the agent is initially unsure about the "type" of his opponent. After playing a bargaining game for a number of rounds, the agent may form a belief about his opponent's type and fine-tune his behavior accordingly. Third, in automated negotiation settings (where agents bargain on behalf of their owners) an agent might need to learn the preferences of his owner first. In this report attention is focussed on the first two kinds of learning.

This section is organized as follows. First, several AI learning techniques are introduced briefly in Section 3.1. Next, Section 3.2 shows how efficient bargaining strategies can be obtained using evolutionary algorithms (EAs). Section 3.3 approaches learning during the negotiation process using Bayesian beliefs. Section 3.4 considers an alternative approach by viewing negotiations as a constraint satisfaction problem. Here, the emphasis lies more on finding acceptable rather than optimal solutions.

3.1 Learning methods

Several learning methods developed in the field of AI are introduced in this section: decision trees, Q-learning, evolutionary techniques and Bayesian beliefs. This list is by no means exhaustive but gives an good impression of the different types of methods that can be relevant in this context.

Two important kinds of learning methods are studied in more detail in subsequent sections. Evolutionary approaches are discussed in Section 3.2. Bayesian beliefs are considered further in Section 3.3.

Decision trees Decision tree induction is one of the simplest supervised learning algorithms [38, Ch. 18]. The inputs of a decision tree are values for a set of attributes (such as the current income of a customer in a bank). The output of a decision tree is usually an action or a decision, usually in the form of a "yes/no" classification (such as, "grant or deny a loan"). Each node in the tree contains an attribute, and the arcs represent different ranges of the value domain for the attribute. Learning is done by adapting the tree in such a way that it remains consistent with the examples presented. The examples contain both the input and the desired output. When a negative example is presented (i.e., when the desired output is not obtained) a node is split and a new attribute is added to the tree. Due to noise negative examples could remain even if all attributes are already used. A simple solution is then to use a majority vote.

In general, a set of training examples can agree with many different decision trees, having a different degree of complexity. Using "Ockham's razor", smaller (i.e., simpler) decision trees are often preferred over larger ones. Unfortunately, the problem of finding the smallest decision tree is an intractable problem [38, p. 535].

Many different learning algorithms have been proposed for solving problems such as efficiency, noise and overtraining. Examples include ID3, C4.5, EPAM, CLS and genetic programming. For an overview see [38, pp. 559-560].

Q-Learning An agent receives feedback each time it performs an action in case of supervised learning. However, in many practical cases feedback is only received at the end of a (long) sequence of actions. A good example is a game like chess: only at the end of play the players know with certainty how well their strategy performs. In learning models like Q-learning, agents also try to evaluate the effect of intermediate actions. Q-learning is a reinforcement learning algorithm [38, p. 528] which learns an action-value function yielding the *expected* utility of a given action in a given state [38, p. 599].

This algorithm maintains a list of so-called Q-values $Q(a, i)$, which denote the expected utility of performing an action a at state i. The action which maximizes the expected utility is selected, and the system moves to a new state j. The Q-value is then updated depending on the Q-value of the new state and the received reward (if available). The following equation can be used [38, p. 613] for updating the Q-value in case of a transition from state i to j by taking action a:

$$Q(a, i) \leftarrow Q(a, i) + \alpha(R(i) + \max_{a'} Q(a', j) - Q(a, i)), \tag{3.1}$$

where $R(i)$ is the actual reward received in state i and α is the learning rate. The value $\max_{a'} Q(a', j)$ represents the expected utility of state j. For example, if the current state i has a relatively low expected utility and the next state j has a high expected utility, the Q-value $Q(a, i)$ is updated in such a way that the difference between these states is reduced. In this way rewards which are given at the terminal state are passed to the other states in the sequence.

As we mentioned before, selecting an action in the current state depends on the expected utility of each action. Hence, a trade-off needs to be made between "exploitation" and "exploration". In other words, should an action be chosen which has already proven itself or do we prefer to try out new actions which might produce even better results? This question of finding an optimal exploration policy has been studied extensively in the subfield of statistical theory that deals with so-called "bandit" problems [38, pp. 610-611]. An application of Q-learning techniques is given in [27].

Q-learning is closely related to learning in classifier systems [14]. Classifiers are rules, which, once activated, activate other rules creating a chain of activations. The last rule in the sequence receives an external reward, which is backpropagated (using a so-called "bucket-brigade" algorithm) to all rules which caused its activation. In addition, new rules are created replacing poorly performing rules (i.e., which generate low rewards). Exploration of new rules is done using a "genetic" algorithm. This learning method is introduced below.

Evolutionary algorithms Evolutionary algorithms (EAs) apply the principles of natural evolution, first discovered by Darwin and Mendel, in a computational setting. The cornerstones of evolution in nature are "survival of the fittest" together with the transfer (with some variation) of genetic material from parents to their children. Transfer of genetic material (DNA) from parents to offspring typically occurs in two steps. During the *recombination* phase the parental chromosomes are paired two-by-two and "crossed over". Errors in this recombination process or external factors like radiation or chemical processes can lead to additional *mutations* of the chromosomes.

The survival and future reproduction of offspring is depending on their "fitness", that is, their ability to gather scarce resources. This process of evolution causes good traits to remain in the population and bad traits to die out in the long run.

Evolutionary algorithms mimic some aspects of these biological processes in a computer [14, 21, 2]. EAs typically use a population of individuals. The individuals are not living organisms in this case, but for instance solutions for a optimization problem or strategies of agents playing a game. These solutions are encoded on a "chromosome", most often consisting of a sequence of binary or real-coded numbers. As in natural ecosystems, the survival of these individuals depends on their fitness. A suitable fitness measure in artificial ecosystems depends on the problem domain. It can for instance be an objective function in case of an optimization problem, or the mean utility obtained by a strategy in a game.

Genetic algorithms (GAs) are a special class of evolutionary algorithms, first developed by Holland [14]. Here, the chromosomes of the individuals are encoded using bit strings. The genetic crossover and mutation operators are used to create new individuals not yet present in the population. Additionally, a selection operator is used to select the fittest (i.e., closest to the optimal solution) individuals which are then allowed to produce offspring. Other classes of evolutionary algorithms include genetic programming, evolution strategies, and evolutionary programming [2].

Bayesian beliefs The meaning of the general term "belief" is depending on the problem domain. In a multi-agent context, beliefs may for instance represent contingent statements (i.e., they could be incorrect) about an agent's environment. To avoid confusion at this point, we therefore continue with a discussion of "Bayesian" beliefs, defined as in the field of probabilistic reasoning.

Bayesian beliefs are used to model an agent's (probabilistic) knowledge of an uncertain environment. Suppose the agent has some a priori knowledge about the likelihood of a set of hypotheses H_i, with $i = 1, ..., n$. Furthermore, the agent has some conditional knowledge about the probability that an event e will occur, given that one of the hypotheses is true. If event e then occurs, the beliefs about the hypotheses are updated using the Bayesian update rule [50]:

$$P(H_i|e) = \frac{P(H_i)P(e|H_i)}{\sum_{k=1}^{n} P(e|H_k)P(H_k)}, \tag{3.2}$$

where $P(H_i|e)$ is the a posteriori probability of H_i and $P(H_i)$ the a priori probability. $P(e|H_i)$ is the conditional probability that event e occurs given hypothesis H_i.

Other techniques Other learning algorithms include classifier systems, neural networks and cellular automata. For a short overview on these techniques, see [9, pp. 13-21]. Moreover, there are mixed approaches, e.g., evolving decision trees using genetic programming and evolving a classifier system using GAs.

3.2 The evolutionary approach

Most of today's automated negotiation systems for e-commerce on the Internet use simple and static negotiation rules. Examples are Kasbah[14] and AuctionBot[15].[16] These examples show that at the moment few systems use techniques from the field of machine learning. Below we will discuss some key papers which address the important question of how to make negotiation systems adaptive. The basic technique used in these papers is the evolutionary approach.

Oliver [28] was the first to demonstrate that a system of adaptive agents can learn effective negotiation strategies. Computer simulations of both distributive (i.e., single issue) and integrative (i.e., multiple issue) "alternating-offers" negotiations are presented in [28]. Binary coded strings represent the agents' strategies. Two parameters are encoded for each negotiation round: a threshold which determines whether an offer should be accepted or not and a counter offer in case the opponent's offer is rejected (and the deadline has not yet been reached). These elementary strategies were then updated in successive generations by a genetic algorithm (GA). A similar model has been investigated in [12, 43]

More elaborate strategy representations are proposed in [20]. Offers and counter offers are generated in this model by a linear combination of simple bargaining tactics (time-dependent, resource-dependent, or behavior-dependent tactics). As in [28], the parameters of these different negotiation tactics and their relative importance weightings are encoded in a string of numbers. Competitions were then held between two separate populations of agents, which were simultaneously evolved by a GA.

Dworman et. al [10] studied negotiations between three players. If two players decide to form a coalition, a surplus is created which needs to be divided among them. The third party gets nothing. Of course, all three players want to be part of the coalition in this case. Moreover, they also want to receive the largest share of the bargaining surplus. Genetic programming was used in this paper to adapt the offers and to decide whether to form a coalition or not. A comparison with game theoretic predictions and human experiments was made.

3.3 Using Bayesian beliefs

When agents have incomplete information about one another, it becomes important to learn about the other agent by observing his behaviour during the negotiation process. Bayesian beliefs are often used to make assumptions about the opponent such as his "type" [17] or his reservation price [49],[50]. These beliefs are updated depending on the opponent's moves.

[14] http://kasbah.media.mit.edu/.
[15] http://auction.eecs.umich.edu/.
[16] An overview of agent-based e-commerce applications is given in [19, 13].

However, once both agents use beliefs to determine their strategies, they also need beliefs about their opponent's beliefs, and so on. This is known as the problem of "outguessing regress" [50]. In game theory this problem is solved by having a limited number of different types of players. The beliefs and preferences of each type are common knowledge, but there is uncertainty about which player is of which type. This theory, suggested by Harsanyi, is a technique for transforming a game of incomplete information into a game of imperfect (but complete) information (see [3], pp. 501-510). In reality however, the number of different types is usually very large, and, moreover, it is not always realistic to assume that the preferences and beliefs of the different types are common knowledge. In more practical applications (such as [17] and [49]), the problem of outguessing regress is circumvented by assuming bounded rationality.

3.4 Distributed constraint satisfaction problem solving

Instead of quantatively encoding the preferences of the negotiating parties one might also use a more qualitative model. Guttman and Maes [13] propose the use of distributed constraint satisfaction problem solving techniques (DCSPs) as an alternative approach to the quantitative multi-attribute utility theory. Concepts like reservation prices can be modeled within this framework using hard constraints. Additionally, this model can in principle incorporate "soft" constraints, which model other preferences such as inter-issue relationships (e.g. "availability is more important to me than a low price"). In the latter case, not all constraints need to be satisfied [13]. Below, DCSP techniques are briefly described. In this "classical" version, only hard constraints are taken into account. Then, an application of these techniques in an e-commerce context is discussed. Finally, we discuss a closely related technique which uses argumentation to resolve conflicts.

Principles of DCSPs Each agent is associated with a variable and a set of domain values for that variable in DCSPs, as well as a set of constraints for certain combination of values. For example, in an e-commerce context a variable might simply be the price, and the reservation values of the agents form constraints on this variable. The goal is then to find assignments of values to all the variables such that all constraints are satisfied [48]. Communication between agents is done by sending messages. An agent can propose an assignment, and the other agent can reply by either confirming or sending a "no good" message. In case of a no good message the set of assignments which violate one or more constraints is also communicated. This mechanism allows asynchronous activities of the agents, that is, there is no need for a central control mechanism.

Negotiations involving multiple parties This approach was used by Oliveira and Rocha [27] for the formation of virtual organizations in an e-commerce environment. The idea is that in order to satisfy some user's need, often a combination of services is needed, which is provided by different companies. The agent representing the user (called the "market agent") negotiates with several organization agents, after which a selection of these organizations is made and a virtual organization is created. During the negotiations process, the bilateral constraints between the market agent and the organizations need to be resolved. After the selection processes, remaining inter-organizational constraints must be resolved. A solution is obtained in this phase by interaction between the selected organizations, using a distributed constraint based algorithm and without any interference of the market agent.

The protocol used during the negotiation phase is as follows. First, each participating organization generates a bid, based on previous experience, and sends this bid to the market agent. A Q-learning technique is then used to determine which bid to make. The actions (i.e., the bids) made are then evaluated using the feedback given by the market agent. The market agent compares the bids using a multi-criteria evaluation method based on qualitative measures (in which only the preference ordering is assumed to be important). The market agent selects the organization which either proposes a satisfactory evaluation, or he chooses the highest evaluation when a deadline is reached. Organizations not selected are given feedback as to which attributes were not satisfactory (i.e., which constraints were not resolved). Negotiations take several rounds, and each round an organization is selected.

Although the basic idea seems to be promising, there are some hooks and eyes to this approach. First, as pointed out in [13], the approach assumes "cooperation" in the sense that the agents do not have any incentive to hide information. The agents negotiate over several attributes, thereby creating a mutually beneficially outcome. This should be a sufficient incentive for the agents to reveal the necessary information. However, we believe that in several cases the individual agents could gain by lying (about for example their valuation of various offers). In particular, this may be the case in the constraint process solving phase. In this phase, an inter-personal utility comparison is made in order to select the best global solution.

Another problem is the evaluation of qualitative constraints. In [27], the constraints are transformed into a quantitative evaluation function, which is needed for the comparison of the various offers. This transformation seems rather arbitrary. Moreover, it remains unclear whether a distinction is made between soft and hard constraints.

Argumentation-based conflict resolution When negotiations involve several issues and the players differ in their evaluations of the issues, a mutually beneficially situation can be achieved and efficiency comes to play an important role in the negotiation process (as described above). However, when agents have incomplete information about each others' preferences negotiations often result in inefficient deals (see Section 2.4). This problem can be resolved using argumentation. This approach resembles the communication between agents in a DCSP setting. The idea is that the agents are able to provide meta-information on why they have a particular objection to a proposal. This way information is exchanged, but without fully disclosing each others' preferences.

A negotiation architecture using this kind of meta-information is described in [30]. This approach was also used in MIT's Tête-à-Tête system[17], a bilateral integrative negotiation system for online shopping [19]. Agents within this framework can: (1) make a new proposal, (2) accept the proposal of the counter agent, (3) criticize a proposal or (4) withdraw from the negotiations. This system uses the notion of a "critique" to enable agents to criticize a particular proposal. A critique is a comment of an agent specifying which part of the proposal he dislikes. In case of a new proposal or critique, the agent can also send additional information. For instance, a proposal may include conditions under which it holds (e.g., I will provide you with X if you provide me with Y).

4. Concluding remarks

The first part of this report reviewed literature on bargaining from the field of game theory. This overview shows that game theory is a very useful tool to analyze bargaining situations in a mathematical fashion. Such a rigorous analysis is only tractable, however, if many details of human interaction, for instance emotions or irrational behavior, are abstracted away. This may undermine the capability of game-theoretical models to explain or predict human behavior.

This aspect may be less problematic when we consider systems in which *artificial* agents interact with each other, because these agents are often designed to behave (in good approximation) in a rational fashion. Game theory may therefore yield fundamental insights in the design of efficient negotiation protocols for automated trading. Furthermore, given a negotiation protocol and under certain assumptions, optimal strategies can sometimes be derived.

Nevertheless, game-theoretical assumptions like common knowledge and perfect rationality often appear to be too strong in modeling practical situations. The issue of common knowledge has been solved only partially in game theory by introducing a theory for players with "imperfect" information. The development of game-theoretic models for boundedly-rational players is only just starting. Our survey shows that techniques from the field of artificial intelligence are potentially very powerful in situations of incomplete information and boundedly-rational players. Learning techniques developed within the AI community can for instance be used to adapt the agents' behavior in complex

[17]http://ecommerce.media.mit.edu/Tete-a-Tete/.

environments and to construct accurate models of the other agents' preferences.

A variety of learning techniques and state-of-the-art applications have been discussed in this report, some of which seem to be very promising for the use in automated negotiations. Hence, we conclude that combining techniques and ideas from game theory and the field of artificial intelligence opens prospects to create robust, stable and intelligent negotiation systems in the near future.

References

1. M. Bac and H. Raff. Issue-by-issue negotiations: The role of information and time preference. *Games and Economic Behavior*, 13(1):125–134, 1996.
2. T. Bäck, U. Hammel, and H.-P. Schwefel. Evolutionary computation: Comments on the history and current state. *IEEE Trans. on Evol. Comp.*, 1(1):3–17, 1997.
3. K. Binmore. *Fun and Games, A Text on Game Theory.* D.C. Heath and Company, 1992.
4. K. Binmore, M. Piccione, and L. Samuelson. Evolutionary stability in alternating-offers bargaining games. *Journal of Economic Theory*, 80:257–291, 1999.
5. K. Binmore, A. Rubinstein, and A. Wolinsky. The Nash bargaining solution in economic modelling. *Rand Journal of Economics*, 17:176–188, 1986.
6. K. Binmore and N. Vulkan. Applying game theory to automated negotiation. *Netnomics*, 1(1):1–9, 1999.
7. F.M.T. Brazier, F. Cornelissen, R. Gustavsson, C.M. Jonker, O. Lindeberg, B. Polak, and J. Treur. Negotiating for load balancing of electricity use. In M.P. Papazoglou, editor, *Proceedings of the 18th International Conference on Distributed Computing Systems (ICDCS).* IEEE Computer Society Press, 1998.
8. L.-A. Busch and I.J. Horstmann. Signaling via an agenda in multi-issue bargaining with incomplete information. *Economic Theory*, 13(3):561–575, 1999.
9. H. Dawid. *Adaptive Learning by Genetic Algorithms: Analytical Results and Applications to Economic Models.* Lecture Notes in Economics and Mathematical Systems, No. 441. Springer-Verlag, Berlin, 1996.
10. G. Dworman, S.O. Kimbrough, and J.D. Laing. Bargaining by artificial agents in two coalition games: A study in genetic programming for electronic commerce. In *Genetic Programming 1996: Proceedings of the First Annual Conference*, pages 54–63. MIT Press, 1996.
11. C. Fershtman. The importance of the agenda in bargaining. *Games and Economic Behavior*, 2:224–238, 1990.
12. E.H. Gerding, D.D.B. van Bragt, and J.A. La Poutré. Multi-issue negotiation processes by evolutionary simulation. In *Workshop on* Complex Behaviour in Economics, Aix-en-Provence, France, May 2000. forthcoming.

13. R.H. Guttman and P. Maes. Cooperative vs. competitive multi-agent negotiations in retail electronic commerce. In *Proceedings of the 2nd International Workshop on Cooperative Information Agents (CIA'98)*, 1998.

14. J.H. Holland. *Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence.* MIT Press, 1992.

15. E. Kalai and M. Smorodinsky. Other solutions to the Nash bargaining problem. *Econometrica*, 43:513–518, 1975.

16. J. Kennan and R. Wilson. Bargaining with private information. *Journal of Economic Literature*, 31:45–104, 1993.

17. S. Kraus. Beliefs, time and incomplete information in multiple encounter negotiations among autonomous agents. *Annals of Mathematics and Artificial Intelligence*, 20:111–159, 1997.

18. R.D. Luce and H. Raiffa. *Games and Decisions.* Wiley, New York, 1957.

19. P. Maes, R.H. Guttman, and A.G. Moukas. Agents that buy and sell. *Communications of the ACM*, 42(3):81–91, 1999.

20. N. Matos, C. Sierra, and N.R. Jennings. Determining successful negotiation strategies: An evolutionary approach. In *Proceedings of 3rd International Conference on Multi-Agent Systems (ICMAS-98)*, pages 182–189, Paris, France, 1998.

21. M. Mitchell. *An introduction to genetic algorithms.* MIT Press, 1996.

22. A. Muthoo. *Bargaining and Applications.* Cambridge University Press, Cambridge, 1999.

23. R.B. Myerson. Utilitarianism, egalitarianism, and the timing effect in social choice problems. *Econometrica*, 49(4):883–897, 1981.

24. J.F. Nash. The bargaining problem. *Econometrica*, 18:155–162, 1950.

25. J.F. Nash. Equilibrium points in N-person games. *Proc. Nat. Acad. Sci. U.S.A.*, 36:48–49, 1950.

26. J.F. Nash. Non-cooperative games. *Annals of Mathematics*, 54:286–295, 1951.

27. E. Oliveira and A.P. Rocha. Agents advances features for negotiation in electronic commerce and virual organisations formation process. In C. Sierra and F. Dignum, editors, *Book on European perspectives on AMEC.* Springer-Verlag, June 2000. To appear, draft available at http://www.iiia.csic.es/AMEC/BOOK/Book.html.

28. J.R. Oliver. A machine learning approach to automated negotiation and prospects for electronic commerce. *Journal of Management Information Systems*, 13(3):83–112, 1996.

29. M.J. Osborne and A. Rubinstein. *Bargaining and Markets.* Academic Press, 1990.

30. S. Parsons, C. Sierra, and N. Jennings. Agents that reason and negotiate by arguing. *Journal of Logic and Computation*, 8(3):261–292, 1998.

31. C. Ponsati and J. Watson. Multiple-issue bargaining and axiomatic solutions. *International Journal of Game Theory*, 26:501–524, 1997.

32. H. Raiffa. *The Art and Science of Negotiation.* Harvard University Press, Cambridge, 1982.

33. J.S. Rosenschein and G. Zlotkin. *Rules of Encounter.* MIT Press, Cambridge, 1994.

34. A.E. Roth. *Axiomatic Models of Bargaining.* Lecture Notes in Economics and Mathematical Systems, No. 170. Springer-Verlag, Berlin, 1979.

35. A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):155–162, 1982.

36. A. Rubinstein. A bargaining model with incomplete information about time preferences. *Econometrica*, 53(5):1151–1172, 1985.

37. A. Rubinstein. *Modeling Bounded Rationality.* The MIT Press, Cambridge, 1998.

38. S. Russell and P. Norvig. *Artificial Intelligence, A Modern Approach.* Prentice-Hall, 1995.
39. R. Selten. Spieltheoretische behandlung eines oligopolmodels mit nachfragetragheit. *Zeitschrift fur die Gesamte Staatswissenschaft*, 12:301–324 and 667–689, 1965.
40. R. Selten. Re-examination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory*, 4:25–55, 1975.
41. J. Maynard Smith. *Evolution and the Theory of Games.* Cambridge University Press, Cambridge, 1982.
42. I. Ståhl. *Bargaining Theory.* Stockholm School of Economics, Stockholm, 1972.
43. D.D.B. van Bragt, E.H. Gerding, and J.A. La Poutré. Equilibrium selection in evolutionary bargaining models. In *Sixth International Conference on Computing in Economics and Finance (CEF2000)*, Barcelona, 2000. forthcoming.
44. E. van Damme. *Stability and Perfection of Nash Equilibria.* Springer, Berlin, 1991.
45. E. van Damme, R. Selten, and E. Winter. Alternating bid bargaining with a smallest money unit. *Games and Economic Behavior*, 2:188–201, 1990.
46. J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior.* Princeton University Press, Princeton, New Jersey, 1944.
47. G. Weiss, editor. *Multiagent Systems, A Modern Approach to Distributed Artificial Intelligence.* MIT Press, 1999.
48. M. Yokoo, H. Durfee, T. Ishida, and K. Kuwabara. Distributed constraint satisfaction for formalizing distributed problem solving. In *Proceedings of the 12th IEEE International Conference on Distributed Computing Systems*, pages 614–621, 1992.
49. D. Zeng and K. Sycara. How can an agent learn to negotiate. In *Intelligent Agents III: Proceedings of ECAI '96 Workshop (ATAL)*, pages 233–244, 1996.
50. D. Zeng and K. Sycara. Bayesian learning in negotiation. *International Journal of Human Computer Systems*, 48:125–141, 1998.